arXiv:1606.07861v2 [cs.DS] 22 Jan 2017

Tight Algorithms for Vertex Cover with Hard Capacities on Multigraphs and Hypergraphs

Sam Chiu-wai Wong
UC Berkeley
samcwong@berkeley.edu

January 24, 2017

Abstract

In this paper we give a f-approximation algorithm for the minimum unweighted Vertex Cover problem with Hard Capacity constraints (VCHC) on f-hypergraphs. This problem generalizes standard vertex cover for which the best known approximation ratio is also f and cannot be improved assuming the unique game conjecture. Our result is therefore essentially the best possible. This improves over the previous 2.155 (for $f = 2$) and $2f$ approximation algorithms by Cheung, Goemans and Wong (CGW).

At the heart of our approach is to apply iterative rounding to a natural LP relaxation that is slightly different from prior works which used (non-iterative) rounding. Our algorithm is significantly simpler and offers an intuitive explanation why f-approximation can be achieved for VCHC. We also present faster implementations of our method based on iteratively rounding the solution to certain CGW-style covering LPs.

We note that independent of this work, Kao [9] also recently obtained the same result.

1 Introduction

The minimum vertex cover problem is one of the earliest NP-hard problems studied in combinatorial optimization. In its most basic form, given a graph $G = (V, E)$ we are asked to find a subset $U \subseteq V$, called vertex cover, so that every edge $e \in E$ intersects U in at least one of its two endpoints. The objective is to minimize the size of U. Curiously, despite decades of efforts the best known algorithms for this problem are the 2-approximation which can either be done by LP relaxation or a simple greedy procedure. The minimum vertex cover problem lends itself to a natural generalization to f-hyergraphs where an edge $e \in E$ can have as many as f endpoints. It is not a difficult matter to generalize the 2-approximation to f-approximation for this version. The seminal result of Khot showed that these algorithms are in fact optimal assuming the Unique Game Conjecture (UGC) [10].

Chuzhoy and Naor [2] initiated the study of vertex cover with *hard capacity* constraints (VCHC) where we have a capacity of $k_v \geq 0$ for each $v \in V$ and (a copy of) v can cover at most k_v of its incident edges. The objective is still to minimize the size of the vertex cover found. They gave a natural LP relaxation for VCHC from which a 3-approximation is derived via randomized rounding for graphs with *no multiple edges*. Their analysis is based on Chebyshev inequality. Subsequently Gandhi et al. [5] improved this to a tight 2-approximation by using Chernoff in place of Chebyshev with a much more involved analysis. Both of these algorithms fail to work for multigraphs (graphs possibly with multiple edges) or hypergraphs essentially because in such cases the random variables in their analyses become unbounded and standard concentration inequalities do not apply.

Progress had been stagnated until Saha and Khuller gave a $\min\{6f, 65\}$ approximation for VCHC on hypergraphs [13]. Their idea is to apply randomized rounding for random variables at different scales to salvage Chernoff. Partly inspired by their result, Cheung, Goemans and Wong (CGW) surprisingly gave simple deterministic rounding algorithms which achieve significantly better approximation ratios of 2.155 (for graphs) and $2f$ [1]. Their method is to formulate the coverage requirement of randomized rounding, used in all previous works, in terms of another LP and study the property of its extreme point solutions. In other words, their approach is a 2-stage LP rounding procedure which solves the same LP relaxation followed by the "coverage requirement LP".

1.1 Our contribution

We propose a new simple approach to the problem based on iterative rounding without using new LPs. Our algorithm achieves the best possible approximation ratio f and essentially settles its approximability. Our approach is inspired by ideas used in previous works, most notably CGW which considers an extreme point solution to certain covering LPs. In hindsight their method suggested the possibility of a better approximation obtained by iterative rounding, which often exploits the structure of extreme point solutions. We also show that when combined with iterative rounding, CGW approach can be extended to give another f-approximation. Although more contrived, this alternate algorithm may be preferred as it involves iteratively rounding the solution to so-called *covering LP*s, which can be solved faster than general LPs using dedicated algorithms [12].

Authors	approx. ratio (graphs, hypergraphs)	multigraphs and hypergraphs okay?
Chuzhoy, Naor [2]	3, *	no
Gandhi et al. [5]	2, *	no
Saha, Khuller [13]	12, $\min\{6f, 65\}$	yes
Cheung, Goemans, Wong [1]	2.155, $2f$	yes
This paper, [9]	2, f	yes

1.2 Other related works

Prior to the work of Chuzhoy and Naor [2] which initiated the study of vertex cover with *hard* capacity constraints, Guha et al. [7] resolved the problem with *soft* capacity constraints (where a vertex can be used an arbitrary number of times) using a clever primal-dual algorithm. Notably, their result holds even for the weighted setting whereas the hard capacity version is as hard as set cover in the weighted case and an approximation ratio of $O(\log n)$ is optimal [2]. Using dependent randomized rounding, another 2-approximation for the soft capacity version was given by Gandhi et al. [6].

2 Preliminaries

Let $G = (V, E)$ be a multigraph. We write $u \in e$ to indicate that u is an endpoint of edge $e \in E$.

The minimum Vertex Cover problem with Hard Capacity constraints (VCHC) is specified by (V, E, k, m), where

- $G = (V, E)$ is the input multigraph,
- For each $v \in V$, m_v denotes the maximum number of copies of v one can select,
- For each $v \in V$, k_v is the number of incident edges (a copy of) v can cover.

A solution to VCHC consists of $(x, y) = (\{x_v\}_{v \in V}, \{y(e, v)\}_{e \in E, v \in e})$. Here x_v is the number of copies of vertex v selected, and the assignment variable $y(e, v) \in \{0, 1\}$ represents whether edge e is covered by v, for each $e \in E$ and $v \in e$. A solution (x, y) is *feasible* for VCHC if

1. For all $v \in V$: $x_v \in \{0, 1, \cdots, m_v\}$,
2. For all $e \in E$: $\sum_{v \in e} y(e, v) = 1$ (i.e. any edge must be covered by one of its endpoints),
3. For all $v \in V$: $|\{e : y(e, v) = 1\}| \leq k_v x_v$ (i.e. the total number of edges assigned to v does not exceed its total capacity).

The objective of VCHC is to find a feasible solution (x, y) for VCHC that minimizes $\sum_{v \in V} x_v$, the size of the vertex cover. As VCHC generalizes the classical minimum vertex cover problem which is already NP-hard, we provide efficient algorithms for finding good approximate solutions. Our approach is based on rounding a fractional solution to the following **LP1** relaxation, which has

been used extensively in the literature [2, 5, 13, 1].

$$\begin{aligned}
\min \ & \sum_{v \in V} x_v \\
\text{s.t.} \ & \sum_{v \in e} y(e,v) = 1 && \forall e \in E && \text{(1a)} \\
& y(e,v) \le x_v && \forall e \in E, v \in e && \text{(1b)} \\
& \sum_{e \in \delta(v)} y(e,v) \le k_v x_v && \forall v \in V && \text{(1c)} \\
& x_v \le m_v && \forall v \in V && \text{(1d)} \\
& x, y \ge 0 && && \text{(1e)}
\end{aligned}$$

Here x_u denotes the number of copies of u selected and $y(e,v)$ indicates whether e is covered by v. The first constraint says that each $e \in E$ should be covered by one of its endpoints, the second ensures that e can be covered by a vertex selected, and the third is the capacity constraint.

The following lemma shows that, when constructing a feasible solution to VCHC, we only need the integrality of x, and not of y. This follows easily by the integrality of flows in networks with integer capacities. We refer readers to [2, 13] for a proof.

Lemma 1 (Chuzhoy and Naor [2], generalized to hypergraphs by Saha and Khuller [13])**.** *If (x,y) is feasible for LP1, and x is integral, there exists an integral y' such that (x,y') is feasible for LP 1, and y' can be found efficiently by a maximum flow computation.*

In light of this lemma, it suffices to identify a feasible integral solution x with a good approximation guarantee.

3 f-approximation for VCHC on f-hypergraphs

Let (x^*, y^*) be an optimal extreme point solution to LP1, and $U = \{u \in V : x^*_u \ge 1/f\}$. A natural idea used in all previous works is round up $u \in U$ which involves only a factor f blowup, and select judiciously a subset of $W = \{w \in V : 0 < x^*_w < 1/f\}$.

Covering tight edges Our iterative rounding scheme[1] is based on the observation that a *tight* edge $e \in \delta(u)$ with $y(e,u) = x_u \ge 1/f$ can be rounded up while respecting the capacity constraint. This follows from the capacity constraint used in the LP1 where R.H.S. is $k_u x_u$. Therefore we may effectively remove e from the LP by covering e with u and decreasing k_u by 1, and solve the new smaller LP relaxation. A similar argument was used in the (non-iterative) rounding algorithms in [2, 5].

Nevertheless, one complication arises as any $x^*_u \ge 1/f$ can in principle drop below $1/f$ in later iterations of the algorithm and end up not being selected, i.e. $x_u = 0$ in the final solution. In this case covering e by u is not justified. Here we introduce the constraint $1/f \le x_u$ to the rescue. It ensures that any u with $x^*_u \ge 1/f$ will stay above $1/f$ ever after.

Fixing $x^*_u = 1/f$ To further simplify the LP, we observe that any $x^*_u = 1/f$ can be readily rounded up and removed from the LP. In terms of cost this is a good idea as the approximation ratio incurred is exactly f, meaning that we are not being lossy. Moreover, it ensures that any

[1] See e.g. [11] for the background on iterative rounding which was introduced by Jain [8].

edge would always have an endpoint in U (see Lemma 5) which is important when we bound the approximation ratio by exploiting the structure of the extreme point solution in the proof of Lemma 6. The idea of examining an extreme point solution was inspired by [1].

Modified LP relaxation We incorporate these insights into **LP2** below. LP2 resembles the form of LP1 with a few important modifications. The first constraint involves $\bar{y}(e, v)$ which is the coverage of v towards e in the final solution. This is a result of fixing $x_u^* = 1/f$ as discussed above. The third constraint has $k_u - |T_u|$ in place of k_u; here $|T_u|$ is the number of tight edges covered by u so we simply subtract $|T_u|$ from the capacity k_u. The fourth constraint now sums over only non-tight edges as any tight edges have already benn covered by its endpoint in U. Finally a new constraint $x_u \geq 1/f$ is introduced to ensure that $x_u^* \geq 1/f$ cannot drop below $1/f$. The last new constraint $x_w \leq 1/f$ is, strictly speaking, not needed but is included to simplify our exposition.

$$
\begin{aligned}
\min \quad & \sum_{v \in V \setminus D} x_v \\
\text{s.t.} \quad & \sum_{v \in e \setminus D} y(e, v) = 1 - \sum_{v \in e \cap D} \bar{y}(e, v) && \forall e \in E \setminus T && \text{(2a)} \\
& y(e, v) \leq x_v && \forall e \in E \setminus T, v \in e \setminus D && \text{(2b)} \\
& \sum_{e \in \delta(u) \setminus T} y(e, u) \leq (k_u - |T_u|) x_u && \forall u \in U_> && \text{(2c)} \\
& \sum_{e \in \delta(w) \setminus T} y(e, w) \leq k_w x_w && \forall w \in W && \text{(2d)} \\
& 1/f \leq x_u \leq m_u && \forall u \in U_> && \text{(2e)} \\
& 0 \leq x_w \leq 1/f && \forall u \in W && \text{(2f)} \\
& y \geq 0 && && \text{(2g)}
\end{aligned}
$$

Algorithm We adopt the following notations:

- $U = \{u \in V : x_u^* \geq 1/f\}$, $W = \{w \in V : 0 < x_w^* < 1/f\}$, $Z = \{z \in V : x_z^* = 0\}$.
- Further divide U into $U_> = \{u \in U : x_u^* > 1/f\}$ and $U_= = \{u \in U : x_u^* = 1/f\}$.
- $T = \bigcup_{u \in V} T_u$ is a disjoint union of edges $e \in T_u$ covered by $u \in V$. We call T the set of tight edges.
- $D \subseteq V$ is the set of vertex v whose $\bar{x}_v$ has been determined.

Our algorithm is based on performing iterative rounding on LP2. It incrementally builds up a feasible solution $(\bar{x}, \bar{y})$ to VCHC for which $\bar{x}$ is integral.

Algorithm 1: Iterative Rounding Algorithm for VCHC

1 Solve LP1 for an extreme point solution (x^*, y^*). Initially $T_u = \emptyset$ and $D = \emptyset$; initialize $U, U_>, U_=, W, Z$ based on (x^*, y^*).;
2 **repeat**
3 **for** $v \in Z$ **do**
4 set $\bar{x}_v = x^*_v = 0$, $\bar{y}(e, v) = y^*(e, v) = 0$ for $e \in \delta(v)$;
5 $D \longleftarrow D \cup \{v\}$;
6 **end**
7 **if** $y^*(e, u) = x^*_u$ *for some* $u \in U$ *and* $e \in \delta(u) \backslash T$ **then**
8 set $\bar{y}(e, u) = 1$, $\bar{y}(e, v) = 0$ for $v \in e \backslash \{u\}$;
9 $T_u \longleftarrow T_u \cup \{e\}$;
10 **end**
11 **for** $u \in U_=$ **do**
12 set $\bar{x}_u = 1$, $\bar{y}(e, u) = y^*(e, u)$ for $e \in \delta(u) \backslash T$;
13 $D \longleftarrow D \cup \{u\}$;
14 **end**
15 Solve updated LP2 for a new extreme point solution (x^*, y^*). Update $U, U_>, U_=, W, Z$ based on (x^*, y^*).
16 **until** $U_= = Z = \emptyset$ *and* $y^*(e, u) < x^*_u \forall u \in U, e \in \delta(u) \backslash T$;
17 For $v \notin D$, set $\bar{x}_v = \lceil x^*_v \rceil$ and $\bar{y}(e, v) = y^*(e, v)$ for $e \in \delta(v) \backslash T$

Analysis The algorithm works mostly by design. First we demonstrate feasibility.

Lemma 2. *Suppose a vertex u satisfies $x^*_u \geq 1/f$ at some time during the execution of the algorithm. We must then have $x^*_u \geq 1/f$ after so long as $u \notin D$. Moreover, in the final solution $\bar{x}_u \geq 1$.*

Proof. As soon as $x^*_u \geq 1/f$, we either have $u \in U_=$ or $u \in U_>$. In the former case, u is immediately inserted into D (line 13) and $\bar{x}_u = 1$. In the latter case the constraint $x_u \geq 1/f$ ensures $x^*_u \geq 1/f$ ever after. Eventually we either have $u \in U_=$ (so $\bar{x}_u = 1$) or $\bar{x}_u = \lceil x^*_u \rceil \geq \lceil 1/f \rceil \geq 1$. □

Lemma 3. *The final solution $(\bar{x}, \bar{y})$ is feasible with $\bar{x}$ integral.*

Proof. The fact that $\bar{x}$ is integral simply follows from the description of the algorithm.

We first argue that all edges are covered. For tight edges $e \in T$, we must have set $\bar{y}(e, u) = 1$ for some $u \in U$ at some point (line 8) so e is covered by u. For non-tight edges $e \notin T$, when the algorithm terminates we have

$$\sum_{v \in e \backslash D} \bar{y}(e, v) = \sum_{v \in e \backslash D} y^*(e, v) = 1 - \sum_{v \in e \cap D} \bar{y}(e, v)$$

so e is indeed covered.

It remains to argue that the capacity constraint is satisfied. Lines 3-6 and 11-14 are clearly okay. For Lines 7-10, by Lemma 2 any $u \in U$ satisfies $\bar{x}_u \geq 1$ so we may simply set $\bar{y}(e, u) = 1$ and subtract 1 from the capacity k_u in the LP. This is exactly why we have $(k_u - |T_u|)x_u$ in the third type of constraints. □

Now we bound the approximation ratio. Our argument consists of two ingredients. The first is to observe that an edge $e \notin T$ always intersects U as only vertices with $x_v \leq 1/f$ is put into D. The second, which is much more crucial, exploits the structure of an extreme point solution to show that there cannot be too many fractional x^*_w left at the end of the while loop.

Lemma 4. *In line 15 of the algorithm, the old (x^*, y^*) (from before this line but restricted to only the variables appearing in updated LP2) is still feasible for updated LP2.*

Proof. Clear by inspection. □

Lemma 5. *When the algorithm terminates, for any $e \notin T$ we have $\sum_{v \in e \setminus D} x^*_v \geq \sum_{v \in e \setminus D} y^*(e,v) \geq |e \setminus D|/f$.*

Proof. Note that vertices are assigned to D in Lines 3-6 and 11-14, where we have $x^*_v \in \{0, 1/f\}$ (note that this x^*_v is the one from that particular iteration of the algorithm and not necessarily the final one). Therefore $\bar{y}(e,v) = y^*(e,v) \leq x^*_v \leq 1/f$, which implies

$$\sum_{v \in e \setminus D} x^*_v \geq \sum_{v \in e \setminus D} y^*(e,v) = 1 - \sum_{v \in e \cap D} \bar{y}(e,v) \geq 1 - |e \cap D|/f \geq |e \setminus D|/f,$$

where the last inequality follows from $|e| \leq f$. □

Now we show that there cannot be too many elements in W by a simple counting argument based on examining the extreme point.

Lemma 6. *When the algorithm terminates, $|W| \leq |U^=|$ where $U^= := \{u \in U : x^*_u = m_u\}$.*

Proof. As an extreme point solution, (x^*, y^*) is obtained by setting some of the constraints as equalities. We call these constraints, which form an **invertible** matrix, **tight**. The proof is based on examining the structure and number of these tight constraints.

First note that $0 \leq x_w \leq 1/f$ and $1/f \leq x_u$ cannot be tight since $U_= = Z = \emptyset$. Similarly, $y(e,u) \leq x_u$ for $u \in U$ cannot be tight since no more edges can be added to T. Furthermore, we disregard any edge $e \subseteq D$ since all of its $y(e,v)$ has been determined.

Observe that by the last lemma each edge $e \notin T$ (with $e \setminus D \neq \emptyset$) must have an $x^*_v \geq y^*(e,v) \geq 1/f$. In other words, $1 + |e \cap W| \leq |e \setminus D|$.

Below we count the number of tight constraints of different types. The first, second and third in the table are self-explanatory.

For the fourth one in the table there is a total of $\sum_{e \notin T} |e \cap W| + |W|$ such constraints but we claim that only $\sum_{e \notin T} |e \cap W|$ can be tight. This follows from the fact that for each $w \in W$, setting all of the $1 + |\delta(w) \setminus T|$ corresponding constraints tight would give a singular system. More precisely, this would give $x_w = y(e,w) = 0$ which renders setting the constraint $\sum_{e \in \delta(w) \setminus T} y(e,w) \leq k_w x_w$ tight unnecessary. In other words, including all of these constraints would give rise to a singular matrix. Thus the total number is at most $\sum_w |\delta(w) \setminus T| = \sum_{e \notin T} |e \cap W|$.

For the fifth one, by Lemma 5 each edge $e \notin T$ (with $e \setminus D \neq \emptyset$) must satisfy $x^*_v \geq y^*(e,v) \geq 1/f$ for some $v \in e$. Therefore at least one of the $|e \cap U_>|$ constraints $y(e,u) \geq 0$ is not tight.

constraints	#tight ones
$\sum_{v \in e \setminus D} y(e,v) = 1 - \sum_{v \in e \cap D} \bar{y}(e,v)$	$= \lvert E \setminus T \rvert$
$\sum_{e \in \delta(u) \setminus T} y(e,u) \leq (k_u - \lvert T_u \rvert) x_u$	$\leq \lvert U_> \rvert$
$x_u \leq m_u$	$= \lvert U^= \rvert$
$y(e,w) \leq x_w$, $y(e,w) \geq 0$ and $\sum_{e \in \delta(w) \setminus T} y(e,w) \leq k_w x_w$ $(w \in W)$	$\leq \sum_{e \notin T} \lvert e \cap W \rvert$
$y(e,u) \geq 0$ $(u \in U_>)$	$\leq \sum_{e \notin T} \lvert e \cap U_> \rvert - 1$

Now the number of tight constraints is at most

$$\begin{aligned}|E\backslash T| + |U_>| + |U^=| + \sum_{e\notin T}(|e\cap W| + |e\cap U_>| - 1) &= |E\backslash T| + |U_>| + |U^=| + \sum_{e\notin T}(|e\backslash D| - 1)\\ &= \sum_{e\notin T}|e\backslash D| + |U_>| + |U^=|.\end{aligned}$$

On the other hand, the number of variables is $|U_>| + |W| + \sum_{e\notin T}|e\backslash D|$. Since there is an equal number of tight constraints and variables, we have $|W| \le |U^=|$. □

We are ready to derive our main theorem.

Theorem 7. *Our algorithm is a f-approximation.*

Proof. Feasibility follows from Lemmas 3 and 1. We bound the approximation ratio. By Lemma 4 the old (x^*_{old}, y^*_{old}) is still feasible for the updated LP so the objective value of the new (x^*_{new}, y^*_{new}) is no worse:

$$\text{cost}(x^*_{new}, y^*_{new}) = \sum_{v\in V\backslash D_{new}} x^*_{new,v} \le \sum_{v\in V\backslash D_{new}} x^*_{old,v} = \text{cost}(x^*_{old}, y^*_{old}) - \sum_{v\in U_{old,=}\cup Z} x^*_{old,v}$$

which says that the cost $|U_{old,=}|$ incurred to round up vertex in $U_{old,=}$ can be charged to $\sum_{v\in U_{old,=}\cup Z} x^*_{old,v} = |U_{old,=}|/f$ with a factor f blowup.

For the last (x^*_{last}, y^*_{last}), the cost of rounding up the remaining $\bar{x}_v = \lceil x^*_{last,v}\rceil$ is

$$\begin{aligned}|W| + \sum_{u\in U^=} m_u + \sum_{u\in U\backslash U^=}\lceil x^*_{last,u}\rceil &\le |U^=| + \sum_{u\in U^=} m_u + \sum_{u\in U\backslash U^=}\lceil x^*_{last,u}\rceil\\ &\le f\left(\sum_{u\in U^=} m_u + \sum_{u\in U\backslash U^=} x^*_{last,u}\right)\\ &= f\cdot\text{cost}(x^*_{last}, y^*_{last}),\end{aligned}$$

where we used Lemma 6 and $m_u \ge 1, f \ge 2, x^*_{last,u} \ge 1/f$. This proves the theorem. □

3.1 Implementation using only the original LP

It is possible to implement a similar algorithm without appealing to LP2, which essentially introduces the new constraint $x_u \ge 1/f$. Without them previous $u \in U$ may fall into W and u may not be selected in the final solution, in which case setting $\bar{y}(e,u) = 1$ for $y^*(e,u) = x^*_u$ is not justified as u does not have the capacity to cover e.

The key idea is to continuously move from the old optimum to the new one. More concretely, consider moving along the line from (x^*_{old}, y^*_{old}) to (x^*_{new}, y^*_{new}). We stop whenever $x^*_v = 1/f$ at an intermediate point (x^*, y^*), where we perform lines 11-14 by fixing $x_v = 1$ and $y(e,u) = y^*(e,u)$, and solve the updated LP again. If we arrive at (x^*_{new}, y^*_{new}) without stopping, then we perform lines 7-10 by covering any tight edge $y^*(e,u) = x^*_u \ge 1/f$ using u, i.e. removing e from the LP and decreasing k_u by 1; and solve the updated LP again (one can also eliminate Z which, as with the previous approach, is only introduced to simplify the exposition). If none of these operations are possible, then we are at an extreme point where the same proof would show that rounding up the remaining vertices would give a f-approximation.

Readers can easily make the previous proof work for this new implementation. We refrain from giving a proof because in the next section, we would give a faster implementation using this idea on another LP with a full proof.

4 Faster Implementation via Solving Covering LPs

The approach in the previous section essentially redistributes the coverage relation $y(e, v)$ and cost x_v from one iteration to the next. Upon a closer examination, readers may have noticed that the constraint $y(e, w) \le x_w$ for $w \in W$ plays a relatively smaller role. One may wonder if there could a faster implementation without solving LP2 again. We answer this in the affirmative in this section. This alternate approach performs iterative rounding on LP3, which is a packing LP and can be solved faster using various specialized algorithms. The formulation of LP3 has been heavily inspired by a similar construction in [1].

4.1 LP relaxation and algorithm

In this section we use the same notations $U, U_>, U_=, W, Z$ as before. Our covering **LP3** is as follows:

$$\min \quad \sum_{w \in W} x_w + \sum_{u \in U_>} x_u$$

$$\text{s.t.} \quad \sum_{w \in W} M(u,w) x_w + (k_u - |T_u|) x_u \ge \sum_{w \in W} M(u,w) x^*_w + \sum_{e \in \delta(u) \setminus T} y^*(e,u) \quad \forall u \in U_> \tag{3a}$$

$$0 \le x_w \le 1/f \quad \forall w \in W \tag{3b}$$

$$1/f \le x_u \le m_u \quad \forall u \in U_> \tag{3c}$$

where

$$M(u,w) = \sum_{e \in E_u \cap \delta(w)} \frac{y^*(e,w)}{x^*_w}.$$

Redistributing coverage At the high level, LP3 is based on doing book-keeping of how edges $\delta(u)$ incident to $u \in U$ are covered. It attempts to redistribute the coverage by maintaining the proportion of capacity used for different $y(e, w)$'s (for a given w). While an edge e can have more than one endpoint in U, we simply arbitrarily assign e to one such u so that E_u is the collection of edges assigned to u.

For $w \in W$ we distribute its capacity in the same proportion as before (hence step 4(a) and the definition of $M(u, w)$; see also Lemma 10). The coverage $y(e, u)$ for $e \in E_u$ would then be the remaining amount not yet covered. However this amount can become negative if the other endpoints of e in W contribute more than before towards covering e. Similarly, if they contribute much less now $y(e, u)$ can become larger than x_u.

The key idea is to slowly move from an old solution (x^*, y^*) to the new (x^*_{new}, y^*_{new}), and stop whenever any of these desired conditions is about to fail (step 5). At this point we may simplify the solution by appropriately modifying x, y.

The algorithm is as follows. Readers should recognize the resemblance to the previous one, except with the notable difference that we need to ensure $y(e, u) \ge 0$ in addition to $x_u \ge 1/f$ and $y(e, u) \le x_u$ (pardon us for not using the algorithm environment here as it would clutter the longer description).

Iterative Rounding Algorithm using Covering LPs

1. Solve **LP1** for an extreme point solution (x^*, y^*). Initially $T_u = \emptyset$ and $D = \emptyset$; initialize $U, U_>, U_=, W, Z$ based on (x^*, y^*).
2. For $v \in Z$, set $\bar{x}_v = x^*_v = 0$, $\bar{y}(e, v) = y^*(e, v) = 0$ for $e \in \delta(v)$ and $D \longleftarrow D \cup \{v\}$.
3. Partition edges into a disjoint union $E \backslash T = \bigcup_{u \in U} E_u$ by assigning $e \notin T$ to an arbitrary E_u where $u \in e \cap U$.
4. Solve updated **LP3** for an extreme point solution x^*_{new} and set
   (a) $y^*_{new}(e, w) = \frac{y^*(e,w)}{x^*_w} x^*_{new,w}$ for $w \in W, e \in \delta(w) \backslash T$
   (b) $y^*_{new}(e, u) = y^*(e, u) + \sum_{w \in e \cap W} (y^*(e, w) - y^*_{new}(e, w))$ for $u \in U, e \in E_u$
   (c) $y^*_{new}(e, u) = y^*(e, u)$ for $u \in U, e \notin E_u \cup T$
5. Let $x^t = (1-t)x^* + tx^*_{new}$ and $y^t = (1-t)y^* + ty^*_{new}$, where t **continuously** increases from 0 to 1. **Stop** whenever:
   (a) $y^t(e, u) = x^t_u$ for $u \in U$ and $e \in \delta(u) \backslash T$. Set $\bar{y}(e, u) = 1$, $\bar{y}(e, v) = 0$ for $v \in e \backslash \{u\}$ and $T_u \longleftarrow T_u \cup \{e\}$.
   (b) $x^t_v = 1/f$ for some v. Set $\bar{x}_u = 1$, $\bar{y}(e, u) = y^*(e, u)$ for $e \in \delta(u) \backslash T$ and $D \longleftarrow D \cup \{u\}$.
   (c) $y^t(e, u) = 0$ for $u \in U$ and $e \in \delta(u) \backslash T$. Set $\bar{y}(e, u) = 0$ and $e \longleftarrow e \backslash \{u\}$.
6. Update $U, U_>, U_=, W, Z$ based on $(x^*, y^*) \longleftarrow (x^t, y^t)$.
7. **Repeat** steps 2-7 until no more updates are possible (reaching x^*_{new} without stopping).
8. For $v \notin D$, set $\bar{x}_v = \lceil x^*_v \rceil$ and $\bar{y}(e, v) = y^*(e, v)$ for $e \in \delta(v) \backslash T$.

First we show that our algorithm is well-defined. Interestingly, this is analogous to Lemma 5 which is used to establish approximation guarantee instead.

Lemma 8. *We have $\sum_{v \in e \backslash D} x^*_v \geq \sum_{v \in e \backslash D} y^*(e, v) \geq |e \backslash D|/f$. In particular, there is some $u \in e \cap U$ so step 3 of the algorithm is well-defined.*

Proof. Note that vertices are assigned to D in steps 2 and 5(b), where we have $x^*_v \leq 1/f$. Therefore $\bar{y}(e, v) = y^*(e, v) \leq x^*_v \leq 1/f$, which implies

$$\sum_{v \in e \backslash D} x^*_v \geq \sum_{v \in e \backslash D} y^*(e, v) = 1 - \sum_{v \in e \cap D} \bar{y}(e, v) \geq 1 - |e \cap D|/f \geq |e \backslash D|/f,$$

where the last inequality follows from $|e| \leq f$. □

Now we argue for feasibility. As before, any $x^*_u \geq 1/f$ would stay above $1/f$ which is necessary to cover tight edges in step 5(a).

Lemma 9. *Suppose u satisfies $x^*_u \geq 1/f$ at some time during the execution of the algorithm. We must then have $x^*_u \geq 1/f$ after so long as $u \notin D$. Moreover, in the final solution $\bar{x}_u \geq 1$.*

Proof. Any such x^*_u cannot drop below $1/f$ thanks to step 5(b) of the algorithm, and would be rounded up eventually. □

We justify the capacity constraint in the next lemma, which goes hand-in-hand with the way **LP3** is formulated.

Lemma 10. *$(x^*, y^*), (x^*_{new}, y^*_{new}), (x^t, y^t)$ satisfy the capacity constraint $\sum_{e\in\delta(w)\setminus T} y(e,w) \le k_w x_w$ for $w \in W$ and $\sum_{e\in\delta(u)\setminus T} y(e,u) \le (k_u - |T_u|)x_u$ for $u \in U$.*

Proof. We proceed by induction.

For $w \in W$, we have $\sum_{e\in\delta(w)\setminus T} y^*_{new}(e,w) = \sum_{e\in\delta(u)\setminus T} \frac{y^*(e,w)}{x^*_w} x^*_{new,w} \le k_w x^*_{new,w}$ since by the induction hypothesis we have $\sum_{e\in\delta(w)\setminus T} y^*_{new}(e,w) \le k_w x^*_w$. Now (x^t, y^t) (and therefore the new (x^*, y^*)) also satisfy the capacity constraint since it is a convex combination of $(x^*, y^*), (x^*_{new}, y^*_{new})$.

For $u \in U$, we have

$$\begin{aligned}
\sum_{e\in\delta(u)\setminus T} y^*_{new}(e,u) &= \sum_{e\in E_u} \left[y^*(e,u) + \sum_{w\in e\cap W} (y^*(e,w) - y^*_{new}(e,w)) \right] + \sum_{e\notin E_u\cup T} y^*(e,u) \\
&= \sum_{e\in\delta(u)\setminus T} y^*(e,u) + \sum_{w\in e\cap W} \left(y^*(e,w) - \frac{y^*(e,w)}{x^*_w} x^*_{new,w} \right) \\
&\le (k_u - |T_u|)x^*_{new,u}
\end{aligned}$$

where the last inequality follows from the first constraint of **LP3**. Hence (x^*_{new}, y^*_{new}) satisfies the capacity constraint for $u \in U$, and so is (x^t, y^t) since it is a convex combination of $(x^*, y^*), (x^*_{new}, y^*_{new})$.

Finally, for the new (x^*, y^*) note that T_u may have changed in size but this is okay by design since $(k_u - |T_u|)x^t_u$ changes exactly by $(\#\text{new elements in } T_u) \cdot x^t_u$. □

Lemma 11. *The final solution $(\bar{x}, \bar{y})$ is feasible with $\bar{x}$ integral.*

Proof. The fact that $\bar{x}$ is integral simply follows from the description of the algorithm. The capacity constraints are satisfied by Lemma 10 and the fact that any $u \in U$ would be rounded up to $\bar{x}_u \ge 1$ eventually (Lemma 9) and therefore can pay for covering the tight edges T_u.

The other constraints $0 \le x_v \le m_v, y \ge 0$ are guaranteed by step 5 of the algorithm where we stop at (x^t, y^t) before they can be violated. □

Finally we prove the approximation guarantee, which again mostly follows from examining the structure of an extreme point solution.

Lemma 12. *When the algorithm terminates, $|W| \le |U^=|$ where $U^= := \{u \in U : x^*_u = m_u\}$.*

Proof. This is very similar to the previous proof. Note that the number of variables is $|W| + |U|$. On the other hand, $0 \le x_w \le 1/f$ and $x_u \ge 1/f$ cannot be tight. The number of tight constrains $x_u \le m_u$ is $|U^=|$ while there can be at most $|U|$ tight first covering constraints. So $|W| \le |U^=|$. □

Theorem 13. *Our algorithm is a f-approximation.*

Proof. This is very similar to Theorem 7. The cost of rounding up intermediate $x^t_v = 1/f$ incurs a factor of f. Since (x^*, y^*) is feasible by design and Lemma 10, it remains to show that the last step incurs a factor of at most f. The same inequality in the proof of Theorem 7 works. □

5 Even Faster Implementation for Graphs ($f = 2$)

For the case of graphs $f = 2$ we may in fact drop the variables x_u from **LP3** altogether. This is exactly the same covering LP used in CGW's two-stage rounding algorithm. In this section we show that by performing iterative rounding while moving from old to new optima, their algorithm actually achieves a 2-approximation for graphs.

In our opinion, this approach nicely explains why the randomized rounding scheme in [5] would work. The algorithm of [5] gives a 2-approximation for VCHC on simple graphs by essentially rounding **LP2** via a simple randomized rounding and performing some patching work after. However, a drawback of their work is that the analysis involves many pages of calculations and gives little insights into why the algorithm would give a 2-approximation.

In a way, their argument is a probabilistic proof that a 2-approximate solution exists for **LP4**, and our one-page analysis can be viewed as a simple deterministic proof of that claim.

LP relaxation Instead of **LP3**, we use the following simpler **LP4**:

$$\begin{aligned} \min \quad & \sum_{w \in W} x_w \\ \text{s.t.} \quad & \sum_{w \in W} M(u,w) x_w \geq \sum_{w \in W} M(u,w) x^*_w \quad \forall u \in U & \text{(4a)} \\ & 0 \leq x_w \leq 1 \qquad\qquad\qquad\quad \forall w \in W & \text{(4b)} \end{aligned}$$

where

$$M(u,w) = \sum_{e=uw \in E \setminus T} \frac{y^*(e,w)}{x^*_w}.$$

We first give an overview of the algorithm. The algorithm is mostly a specialization of the previous one to $f = 2$ but the constraint $x_u \geq 1/2$ is now redundant as an edge cannot have both endpoints in W for graphs. One crucial difference in the algorithm is that we are rounding up $y^*(e,u)$ whenever $y^*(e,u) \cdot (\lceil x^*_u \rceil / x^*_u) \geq 1$ (rather than just when $x^*_u = y^*(e,u)$). This is still okay as the final capacity is $k_u \lceil x^*_u \rceil$ so we can blow up the coverage by a factor of $\lceil x^*_u \rceil / x^*_u$.

As for the analysis, to bound the cost of the extreme point solution it is not enough to use only a cardinality inequality like $|W| \leq |U^=|$. Instead we need a finer *matching* structure between U and W that was also used by CGW.

Our algorithm is as follows. Unlike before, *we never redefine or update U and W.*

Iterative Rounding Algorithm using Simpler Covering LPs for Graphs

1. Solve **LP1** for x^*, y^*. Let $T = \emptyset$.
2. For any $u \in U$ and $e \in \delta(u) \backslash T$ satisfying $y^*(e,u) \cdot (\lceil x_u^* \rceil / x_u^*) \geq 1$, set $T \longleftarrow T \cup \{e\}$.
3. Solve updated **LP4** for an extreme point optimum x_{new}^*.
4. Consider $x^t = (1-t)x^* + tx_{new}^*$, where t **continuously** increases from 0 to 1. Let
$$y^t(e,w) = \frac{y^*(e,w)}{x_w^*} x_w^t,\ y^t(e,u) = 1 - y^t(e,w) \quad \forall e = uw \in (U \times W) \backslash T$$
5. Stop whenever $y^t(e,u) \cdot (\lceil x_u^* \rceil / x_u^*) = 1$ for some $u \in U$ and $e \in (U \times W) \backslash T$, set $T \longleftarrow T \cup \{e\}$.
6. Update $M(u,w)$ with $(x^*, y^*) \longleftarrow (x^t, y^t)$. Repeat steps 3-6.
7. If $t = 1$ (i.e. no pair (e,u) satisfies the condition in step 5 for any intermediate t), output $\bar{x}_v = \lceil x_v^* \rceil$ with a corresponding y defined by
$$\bar{y}(e,u) = y^*(e,u),\ \bar{y}(e,v) = y^*(e,v) \quad \forall e \notin T$$
$$\bar{y}(e,u) = 1, \bar{y}(e,v) = 0 \quad \forall e = uw \in T, u \in U, w \in W$$

One difference is that we are rounding up $y^*(e,u)$ whenever $y^*(e,u) \cdot (\lceil x_u^* \rceil / x_u^*) \geq 1$ (rather than just when $x_u^* = y^*(e,u)$). This is still okay as the final capacity is $k_u \lceil x_u^* \rceil$ so we can blow up the coverage by a factor of $\lceil x_u^* \rceil / x_u^*$. Curiously this modification is needed to handle the case $1 < x_u^* \leq 2$ when bounding the approximation ratio.

The feasibility of the algorithm is largely the same as before so we skip the proof (in fact, it is more like an easier special case). To bound the cost of the extreme point solution, it is not enough to use only a cardinality inequality like $|W| \leq |U^=|$ (one fact, now one does not have this but $|W| \leq |U|$). Instead we need a finer *matching* structure between U and W as given in Lemma 15.

Lemma 14. *$(\bar{x}, \bar{y})$ is a feasible solution to VCHC.*

Proof. Similar to Lemmas 10 and 11. □

Lemma 15 (Similar to Theorem 3.2 of [1])**.** *For any extreme point solution x^* to LP4, let $W_f = \{w \in W : 0 < x_w^* < 1\}$. Then there exists a matching between W_f and U that fully matches W_f.*

Proof. Fractional x_w^* can only arise from setting $|W_f|$ covering constraints $\sum_{w \in W} M(u,w) x_w \geq \sum_{w \in W} M(u,w) x_w^*$ tight. Consider the system of equations $Ax = b$ obtained from setting these $|W_f|$ constraints tight. Note that A is $|W_f| \times |W_f|$ with the rows and columns indexed by $|U|$ and $|W_f|$ respectively. A is invertible so
$$\det A = \sum_{\text{permutation } \pi} \left(\pm \prod_i a_{i,\pi(i)} \right) \neq 0$$
which implies that there is a permutation with all $a_{i,\pi(i)} \neq 0$. Such a permutation gives our desired matching since nonzero entries of A corresponds to $M(u,w) \neq 0$ in which case uw is an edge. □

Theorem 16. *$(\bar{x}, \bar{y})$ is a 2-approximation.*

Proof. The cost of U and $x_v = 0$ clearly incurs a factor 2 blowup only. We need to account for the cost of selecting W_f. Now that we don't have $|W_f| \leq |U^=|$, we cannot naively charge $|W_f|$ to $U^=$. However, we have a matching between W_f and U and it suffices to show that for any edge $e = uw$ in the matching,

$$\lceil x_u^* \rceil + 1 = \lceil x_u^* \rceil + \lceil x_w^* \rceil \leq 2(x_u^* + x_w^*).$$

Note that we have $1 = y^*(e,u) + y^*(e,w) \leq x_u^* + x_w^*$ and $y^*(e,u) \cdot (\lceil x_u^* \rceil / x_u^*) \leq 1$ (otherwise step 5 would have applied). We have three cases:

- $\lceil x_u^* \rceil = 1$. Then we are done.
- $\lceil x_u^* \rceil \geq 3$. Then $\lceil x_u^* \rceil + 1 \leq 2x_u^*$.
- $x_u^* + x_w^* \geq 1.5$ and $x_u^* \leq 2$. Then $\lceil x_u^* \rceil + 1 \leq 3 \leq 2(x_u^* + x_w^*)$.
- $x_u^* + x_w^* \leq 1.5$ and $1 < x_u^* \leq 2$, which cannot happen since $1 = y^*(e,u) + y^*(e,w) \leq x_u^*/2 + x_w^* = x_u^* + x_w^* - x_u^*/2 < 1.5 - 1/2 = 1$. Contradiction.

□

In hindsight, CGW came close but failed to obtain a 2-approximation because one does not have $x_{new,u}^* + x_{new,w}^* \geq 1$.

6 Open Problem

In this paper we have settled the approximability of VCHC. While it is UGC-hard to do better [10], the current best NP-hardness inapproximability stands at $f - 1$ [3] and 1.36 for graphs [4]. Is there any hope of proving that it is NP-hard to beat f? In principle it should be easier than doing it for vertex cover since VCHC is more general.

Acknowledgement

We thank Wang Chi Cheung, Lap Chi Lau and Christos Papadimitriou for several helpful discussions. We also thank the anonymous reviewers for many helpful comments that improved the presentation of this paper.

References

[1] Wang Chi Cheung, Michel X Goemans, and Sam Chiu-wai Wong. Improved algorithms for vertex cover with hard capacities on multigraphs and hypergraphs. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1714–1726. SIAM, 2014.

[2] Julia Chuzhoy and Joseph Naor. Covering problems with hard capacities. *SIAM J. Comput.*, 36(2):498–515, 2006.

[3] Irit Dinur, Venkatesan Guruswami, Subhash Khot, and Oded Regev. A new multilayered pcp and the hardness of hypergraph vertex cover. *SIAM J. Comput.*, 34(5):1129–1146, 2005.

[4] Irit Dinur and Samuel Safra. On the hardness of approximating minimum vertex cover. *Annals of mathematics*, pages 439–485, 2005.

[5] Rajiv Gandhi, Eran Halperin, Samir Khuller, Guy Kortsarz, and Aravind Srinivasan. An improved approximation algorithm for vertex cover with hard capacities. *J. Comput. Syst. Sci.*, 72(1):16–33, 2006.

[6] Rajiv Gandhi, Samir Khuller, Srinivasan Parthasarathy, and Aravind Srinivasan. Dependent rounding in bipartite graphs. In *FOCS*, pages 323–332, 2002.

[7] Sudipto Guha, Refael Hassin, Samir Khuller, and Einat Or. Capacitated vertex covering with applications. In *SODA*, pages 858–865, 2002.

[8] Kamal Jain. A factor 2 approximation algorithm for the generalized steiner network problem. *Combinatorica*, 21(1):39–60, 2001.

[9] Mong-Jen Kao. Iterative partial rounding for vertex cover with hard capacities. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2638–2653. SIAM, 2017.

[10] Subhash Khot and Oded Regev. Vertex cover might be hard to approximate to within 2-epsilon. *J. Comput. Syst. Sci.*, 74(3):335–349, 2008.

[11] Lap Chi Lau, Ramamoorthi Ravi, and Mohit Singh. *Iterative methods in combinatorial optimization*, volume 46. Cambridge University Press, 2011.

[12] Serge A Plotkin, David B Shmoys, and Éva Tardos. Fast approximation algorithms for fractional packing and covering problems. *Mathematics of Operations Research*, 20(2):257–301, 1995.

[13] Barna Saha and Samir Khuller. Set cover revisited: Hypergraph cover with hard capacities. In *ICALP (1)*, pages 762–773, 2012.